arXiv:1609.05920v1 [math.OC] 19 Sep 2016

Line Search for Generalized Alternating Projections

Mattias Fält and Pontus Giselsson

September 21, 2016

Abstract

This paper is about line search for the generalized alternating projections (GAP) method. This method is a generalization of the von Neumann alternating projections method, where instead of performing alternating projections, relaxed projections are alternated. The method can be interpreted as an averaged iteration of a nonexpansive mapping. Therefore, a recently proposed line search method for such algorithms is applicable to GAP. We evaluate this line search and show situations when the line search can be performed with little additional cost. We also present a variation of the basic line search for GAP—the projected line search. We prove its convergence and show that the line search condition is convex in the step length parameter. We show that almost all convex optimization problems can be solved using this approach and numerical results show superior performance with both the standard and the projected line search, sometimes by several orders of magnitude, compared to the nominal method.

1 Introduction

Alternating projections is a well known method for feasibilty problems, where the objective is to find a point in the intersection of (convex) sets. The method alternates projections onto the sets. It was first introduced for half-spaces [1] and later generalized to more general sets [6]. In practice, the method is often quite slow. A generalization to this method was proposed in [13], which is based on performing relaxed projections onto the sets instead of standard projections. A relaxation parameter defines how far the relaxed projection should go towards or past the projection point. Depending on the relaxation parameters, it can be shown that the method is an averaged iteration of a nonexpansive mapping, where fixed-points to the mapping correspond to solutions to the feasibility problem. Many variations and extensions of this basic method has been proposed and studied, with linear or sublinear convergence estimates [2, 3]. We present a framework for several of these generalizations and collect the relevant results. The framework includes well known methods such as the alternating projections and the generalized Douglas-Rachford algorithm for feasibility problems. These are first order methods that scale better with the number of variables compared to interior point algorithms and usually have a low computational cost per iteration. They are therefore useful for solving large-scale convex optimization

problems, where the computational complexity is too large for other algorithms. The practical rate of convergence can however be slow and is dependent on preconditioning for good performance. A good preconditioning can be hard to find and is usually problem specific.

Line search is a well established concept in optimization and is often used to improve practical performance of a method. Typically, it assumes that a descent direction for the objective function is at hand, and it accepts points with sufficient decrease and possibly some condition on the slope [5, 15]. For averaged iterations of nonexpansive mappings, descent directions are not obtained in general. In the recent paper [11], a line search method that can be applied to averaged iterations was proposed. The line search is performed in the direction of the fixed-point residual, which is the direction obtained by applying the nonexpansive operator. Instead of being based on objective function value decrease, it relies on a decrease in the norm of the fixed-point residual.

The main contribution of this paper is an alternative to the basic line search for GAP—the projected line search. This is developed for the case with two sets, where one is affine, and we show that most convex optimization problems can be posed on this form. The projected line search method performs line search, not in the residual direction, but in its projection on the affine set. We prove that the method converges, and show that the line search condition is convex in the step length parameter. We also present a numerical example that illustrates the properties of the methods, and show that the projected line search can achieve superior performance.

Section 2 contains some background and notation. In Section 3 we present the generalized alternating projections algorithm and collect relevant results. In Section 4 we show how the line search in [11] can be applied to this algorithm. The projected line search is presented in Section 5 together with some basic results. An overview of how this method can be used to solve a large set of convex optimization problems is presented in Section 6, and a numerical example is presented in Section 7.

2 Background and Notation

The notation $\langle \cdot, \cdot \rangle$ is used for scalar product and Id is the identity operator. The fixed-points of an operator T are denoted $\mathrm{fix}T$, i.e. $\mathrm{fix}T = \{x \in \mathbb{R}^n \mid Tx = x\}$, and the fixed-point residual $r(x)$ for a point x is defined as $r(x) := Tx - x$. An operator $T : \mathbb{R}^n \to \mathbb{R}^n$ is said to be nonexpansive if it satisfies $\|Tx - Ty\|_2 \leq \|x - y\|_2$ for all $x, y \in \mathbb{R}^n$, and an operator S is α-averaged, with $\alpha \in (0, 1)$, if it can be written as $S = (1 - \alpha)\mathrm{Id} + \alpha \mathrm{T}$ for some nonexpansive T. Π_C is the orthogonal projection onto the closed, convex and nonempty set C, i.e. $\Pi_C(x) = \arg\min_{y \in C}(\|y - x\|_2)$.

3 Generalized Alternating Projections

Generalized alternating projections is an algorithm for finding a point in the intersection of p sets C_i with $i = 1, \ldots, p$, i.e., to find a point $x \in C_1 \cap \cdots \cap C_p$. Thoughout this paper, we assume that that sets C_i are nonempty, closed and

convex, and that

$$C_1 \cap ... \cap C_p \neq \emptyset, \tag{1}$$

i.e., that a common feasible point exists.

To define the algorithm, we introduce the under ($\alpha \in (0,1)$) and over ($\alpha \in (1,2]$) relaxed projection on the set C as follows:

$$P_C^\alpha = (1-\alpha)\mathrm{Id} + \alpha\Pi_C \tag{2}$$

where $\alpha \in (0,2]$ and Π_C is the orthogonal projection onto the set C. For $\alpha = 1$, we get the standard projection $P_C^1 = \Pi_C$ and for $\alpha = 2$, we get the reflection $P_C^2 = 2\Pi_C - \mathrm{Id} =: \mathrm{R_C}$. The relaxed projector is $\frac{\alpha}{2}$-averaged for $\alpha \in (0,2)$ and nonexpansive for $\alpha = 2$, since Π_C is firmly nonexpansive, see [4, Corollary 4.29, Example 12.25, Proposition 12.27].

The generalized alternating projections method (GAP) is:

$$x^{k+1} = (1-\alpha)x^k + \alpha P_{C_p}^{\alpha_p} P_{C_{p-1}}^{\alpha_{p-1}} \cdots P_{C_1}^{\alpha_1} x^k. \tag{3}$$

For simplicity, we introduce the GAP operator T as

$$T = (1-\alpha)\mathrm{Id} + \alpha \mathrm{P}_{\mathrm{C_p}}^{\alpha_\mathrm{p}} \mathrm{P}_{\mathrm{C_{p-1}}}^{\alpha_{\mathrm{p-1}}} \cdots \mathrm{P}_{\mathrm{C_1}}^{\alpha_1} \tag{4}$$

to arrive at the notationally more convenient iteration $x^{k+1} = Tx^k$ for (3).

The algorithm (3) generalizes the classical alternating projections method, since if $\alpha = \alpha_i = 1$, we get

$$x^{k+1} = \Pi_{C_p} \Pi_{C_{p-1}} \cdots \Pi_1 x^k.$$

For $p = 2$, the generalized Douglas-Rachford algorithm for feasibility problems [9, 14], also falls under the formulation (3) by letting $\alpha_1 = \alpha_2 = 2$. Then

$$x^{k+1} = (1-\alpha)x^k + \alpha R_{C_2} R_{C_1} x^k$$

where $R_C = 2\Pi_C - \mathrm{Id}$ is a reflector. These two algorithms are illustrated for a simple 2-dimensional problem in Figure 1.

Below, we present some basic results on the algorithm (3). Most of these are known but spread out in the literature so we collect them here for convenience of the reader. To this end, we let

$$\beta := \frac{\sum_{i=1}^p \frac{\alpha_i}{2-\alpha_i}}{1 + \sum_{i=1}^p \frac{\alpha_i}{2-\alpha_i}}, \tag{5}$$

and state the following assumptions on α.

Assumption 1 *Suppose that either of the following holds:*

A1. $\alpha \in (0, \frac{1}{\beta})$ *with* β *in* (5) *and that* $\alpha_i \in (0,2)$ *for* $i = 1, \ldots, p$

A2. $\alpha \in (0,1)$ *and* $\alpha_i \in (0,2]$ *for* $i = 1, \ldots, p$ *with at most one* $\alpha_i = 2$

A3. $\alpha \in (0,1)$ *and* $p = 2$ *with* $\alpha_1 = \alpha_2 = 2$.

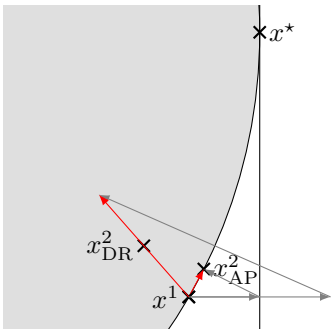

Figure 1: Illustration of the generalized alternating projections for two different settings on a 2-dimensional problem with intersection $x^\star$. The first set is the vertical line, and the second is the shaded area. The point x^2_{AP} is obtained by an alternating projections step ($\alpha_1 = \alpha_2 = 1$, $\alpha = 1$) and x^2_{DR} is obtained by a Douglas-Rachford step ($\alpha_1 = \alpha_2 = 2$, $\alpha = 0.5$). The red arrows represent the residuals $P_{C_2}^{\alpha_2} P_{C_1}^{\alpha_1} x^1 - x^1$ along which we will perform line search in Section 4.

These assumptions imply that the GAP operator T is averaged. This is shown next.

Proposition 1 *Suppose that Assumption 1 with case A1 holds. Then the GAP operator T in (4) is averaged with constant $\alpha\beta \in (0,1)$, with β in (5). Suppose that Assumption 1 with case A2 or A3 holds. Then T is averaged with constant $\alpha \in (0,1)$.*

A proof is found in the Appendix.

Next, we show a result on the fixed-point set of the GAP operator in (4).

Proposition 2 *Suppose that Assumption 1 holds with case A1 or A2 and that $C_1 \cap \cdots \cap C_p \neq \emptyset$. Then $\mathrm{fix} T = C_1 \cap \cdots \cap C_p$, where T is the operator in (4).*

A proof is found in the Appendix.

The main convergence result for the algorithm now follows directly from [4, Theorem 5.14] under assumption A1 or A2 since T is averaged and its fixed-point set is $C_1 \cap \cdots \cap C_p$.

Proposition 3 *Suppose that Assumption 1 holds with case A1 or A2 and that $C_1 \cap \cdots \cap C_p \neq \emptyset$. The fixed-point residuals $r(x^k)$ converge to 0 and the iterates x^k converge to a point in the intersection $C_1 \cap \cdots \cap C_p$, as $k \to \infty$ in algorithm (3).*

Algorithm (3) with case A3 in Assumption 1 corresponds to generalized Douglas-Rachford applied to feasibility problems. The properties in this case are slightly different but well known, and we summarize them below [4, Proposition 25.1, Theorem 25.6].

Proposition 4 *Suppose that Assumption 1 holds with case A3 and that $C_1 \cap C_2 \neq \emptyset$, then the fixed-point set satisfies $\Pi_{C_1} \mathrm{fix} T = C_1 \cap C_2$. Additionally, the fixed-point residuals $r(x^k)$ in algorithm (3) converge to 0 as $k \to \infty$ and the iterates x^k converge to a point x such that $\Pi_{C_1} x \in C_1 \cap C_2$.*

We see that we need to monitor the sequence $\Pi_{C_1} x^k$ to find a feasible point in the Douglas-Rachford case. For other choices of α_i, it is also typically better to monitor (one of) the projected sequences than x^k in (3) to faster find an intersection point (up to numerical accuracy).

4 Line search

A method for applying line search on algorithms based on iterating averaged operators was recently proposed in [11]. The method was shown to often improve practical convergence. In this section we describe how the method can be applied to generalized alternating projections. We also repeat the result that under some assumptions on the sets, the line search can be carried out with little additional cost compared to a basic iteration.

The line search algorithm can be applied to averaged iterations of the form

$$x^{k+1} = (1-\alpha)x^k + \alpha S x^k = x^k + \alpha(Sx^k - x^k), \tag{6}$$

where $\alpha \in (0,1)$ and $S : \mathbb{R}^n \to \mathbb{R}^n$ is nonexpansive. GAP is precisely on this form with $S = P_{C_p}^{\alpha_p} P_{C_{p-1}}^{\alpha_{p-1}} \cdots P_{C_1}^{\alpha_1}$. The second expression in (6) shows that an averaged iteration performs a step with length α in the residual direction $r(x) = Sx - x$. We call this the nominal step $\bar{x}^k := x^k + \alpha r(x^k)$. The residual direction is illustrated in Figure 1.

The line search scheme presented in [11], suggests to perform line search in the residual direction. To do this, the α that multiplies the residual direction should be chosen on-line. The algorithm with line search can be written as:

$$x^{k+1} := x^k + \alpha_k(Sx^k - x^k) := x^k + \alpha_k r(x^k) \tag{7}$$

where the line search parameter α_k must satisfy either $\alpha_k = \alpha$, i.e., we take the nominal step $\bar{x}^k$, or $\alpha_k \in (\alpha, \alpha^{\max}]$ is such that

$$\|r(x^{k+1})\|_2 \leq (1-\epsilon)\|r(\bar{x}^k)\|_2 \tag{8}$$

where $\epsilon \in (0,1)$ and $\alpha^{\max} \geq \alpha$ are fixed algorithm parameters. To accept a step length α_k in the line search, the residual $r(x)$ should be smaller for the next iterate x^{k+1} than for the nominal step $\bar{x}^k$. This preserves the non-increasing property of the fixed-point residual $\|r(x^{k+1})\|$, even when line search is used. As shown in [11], this is enough to, e.g., guarantee convergence of the residual sequence. An appropriate α_k can for example be selected using a simple forward or backward tracking.

The following form of the algorithm shows which computations are needed in each iteration:

$$r^k := Sx^k - x^k \tag{9}$$

$$\bar{x}^k := x^k + \alpha r^k \tag{10}$$

$$\bar{r}^k := S\bar{x}^k - \bar{x}^k \tag{11}$$

$$x^{k+1} := x^k + \alpha_k r^k, \tag{12}$$

where $S = P_{C_p}^{\alpha_p} P_{C_{p-1}}^{\alpha_{p-1}} \cdots P_{C_1}^{\alpha_1}$. The criterion for line search, i.e. accepting $\alpha_k \neq \alpha$ in (12), can be written

$$\|r^{k+1}\|_2 = \|Sx^{k+1} - x^{k+1}\|_2 \leq (1-\epsilon)\|\bar{r}^k\|_2 \tag{13}$$

where $x^{k+1} = x^k + \alpha_k r^k$ as seen in (12). This general form of the algorithm reveals that we need to compute $S(x^k + \alpha_k r^k)$ for each candidate α_k to verify (8), as well as calculating $S\bar{x}^k$ each iteration. So, to evaluate a candidate point in the line search is roughly as costly as performing one basic step in the algorithm. This may or may not be too costly compared to what is saved due to the line search.

In the following, we will show that sometimes many candidate points can be evaluated in the line search with little additional cost. In the case where the sets $C_n, \ldots, C_1$ are affine, i.e. $C_i = \{x \in \mathbb{R} \mid A_i x = b_i\}$, the projection $z = \Pi_{C_i} x$ is affine and given by the solution to the KKT conditions of the projection:

$$\begin{bmatrix} I & A_i^T \\ A_i & 0 \end{bmatrix} \begin{bmatrix} z \\ \lambda \end{bmatrix} = \begin{bmatrix} x \\ b_i \end{bmatrix}.$$

The relaxed projections will therefore also be affine:

$$P_{C_i}^{\alpha_i} x = (1 - \alpha_i)x + \alpha_i \begin{bmatrix} I & 0 \end{bmatrix} \begin{bmatrix} I & A_i^T \\ A_i & 0 \end{bmatrix}^{-1} \begin{bmatrix} x \\ b_i \end{bmatrix}.$$

It follows that the composition $P_{C_n}^{\alpha_n} \cdots P_{C_1}^{\alpha_1}$ is affine, and GAP (3) can be written as:

$$x^{k+1} = (1 - \alpha)x^k + \alpha S_2 S_1$$

where $S_1 x = P_{C_n}^{\alpha_n} \cdots P_{C_1}^{\alpha_1} x =: Fx + h$, with F and h implicitly defined, and $S_2 = P_{C_p}^{\alpha_p} \cdots P_{C_{n+1}}^{\alpha_{n+1}}$. The following iterations show that several candidate α_k can be tested, without multiple evaluations of S_1 [11]:

$$r^k := S_2(Fx^k + h) - x^k \tag{14}$$

$$\bar{x}^k := x^k + \alpha r^k \tag{15}$$

$$\bar{r}^k := S_2\left(Fx^k + h + \alpha F r^k\right) - \bar{x}^k \tag{16}$$

$$x^{k+1} := x^k + \alpha_k r^k \tag{17}$$

where α_k is selected so that

$$\|S_2(Fx^{k+1} + h) - x^{k+1}\|_2 \leq (1 - \epsilon)\|\bar{r}^k\|_2. \tag{18}$$

The computed quantity $Fx^{k+1} = Fx^k + \alpha_k F r^k$ is then reused in (14), (16) and (18) in the following iteration. Therefore, we only need to compute Fx^0 and Fr^k for all k to evaluate any number of candidate α_k in any number of line searches. If the cost of applying S_2 is negligible, then the line search will result in no significant increase in computation per iteration.

5 Projected line search

In this section we present an alternative to the standard line search, that we call projected line search. We present this line search for feasibility problems with two sets, $C_1 = C$ and $C_2 = D$, where C is affine. This method does not select the next iterate in the direction of the residual but rather along its projection on the affine set.

The proposed algorithm, with $S = P_D^{\alpha_2} P_C^{\alpha_1}$, is:

$$r^k := Sx^k - x^k \tag{19}$$

$$\bar{x}^k := x^k + \alpha r^k \tag{20}$$

$$\bar{r}^k := S\bar{x}^k - \bar{x}^k \tag{21}$$

where the next step is to either take a nominal step:

$$x^{k+1} := x^k + \alpha r^k = \bar{x}^k \tag{22a}$$

or line search is performed:

$$x^{k+1} := \Pi_C(x^k + \alpha_k r^k). \tag{22b}$$

To accept the line search in (22b), the line search parameter $\alpha_k \in (\alpha, \alpha^{\max}]$ must satisfy the following constraint, where i_{LS} is the index when the last line search was performed, and $r(x^{i_{\text{LS}+1}}) = Sx^{i_{\text{LS}+1}} - x^{i_{\text{LS}+1}}$ is the residual at the following iteration:

$$\|r(x^{k+1})\|_2 \leq (1-\epsilon)\|r(x^{i_{\text{LS}+1}})\|_2. \tag{23}$$

Compared to the algorithm with basic line search method, the difference is that the candidate points in the projected line search are projected onto the set C. The test for accepting a line search is also different. Instead of comparing the norm of the next residual $r(x^{k+1})$ to the residual of the nominal step $r(\bar{x}^k)$, we compare it to the residual in the last step that was chosen by a line search, $r(x^{i_{\text{LS}+1}})$. The reason for comparing the residual to the iterate $x^{i_{\text{LS}+1}}$ is that the projected line search often increases the residual compared to the nominal step $\bar{x}^k$. However, by ensuring that the residual $r(x^{k+1})$ is smaller than $r(x^{i_{\text{LS}+1}})$, we can guarantee that it will eventually decrease. This is proven for general line search schemes in [11] and we state it for the projected line search below.

Theorem 1 *Assume that Assumption 1 holds and the projected line search algorithm* (19)-(22b) *is used with line search criteria* (23)*. Then the fixed-point residual* $r(x^k) = Sx^k - x^k$ *will converge to* 0 *as* $k \to \infty$*.*

We now show two additional properties of the projected line search.

Theorem 2 *Assume that the set* C *is affine, then the projected line search condition* (23) *is convex in the step length* α_k *and the norm of the residual simplifies to* $\|r(x^{k+1})\|_2 = \alpha_2 \text{dist}_D(x^{k+1})$*.*

A proof is found in the Appendix.

This result implies that we are not restricted to forward or back tracking schemes when finding step length α_k. We can perform, e.g., golden section search on the line search condition, or bisection on its gradient. This way the number of candidate points to be evaluated, before a good/optimal point is found, can be reduced. The theorem also illustrates that minimizing the left hand side of the line search condition is equivalent to minimizing the distance between the two sets along the line search direction. This gives an intuitive explanation to why this is a reasonable objective function.

We showed that the standard line search could be performed without significant extra computational cost in some cases, the same is true for the projected line search. If C is affine, then $x^{k+1} = \Pi_C(x^k + \alpha_k r^k)$ can be evaluated for several α_k without any significant cost since the linear parts of $\Pi_C x^k$ and $\Pi_C r^k$ are known from previous steps, in the same way as for the basic line search. From Theorem 2 we know that evaluating the residual simplifies to evaluating the distance from x^{k+1} the set D. Thus, if C is affine and D is relatively cheap to project on, the line search will incur no significant cost.

6 Cone programming

Many convex optimization problems, including LP, QP, SOCP, SDP, and in particular problems that can be solved using optimization modeling interfaces such as CVX [12], CVXPY [8], Convex.jl [17], can be written as cone programs of the form

$$\begin{array}{ll} \text{minimize} & c^T x \\ \text{subject to} & Ax + s = b \\ & s \in \mathcal{K} \end{array}$$

where $\mathcal{K}$ is a (product of) nonempty, closed and convex cones. The dual of this problem is

$$\begin{array}{ll} \text{maximize} & -b^T y \\ \text{subject to} & -A^T y = c \\ & y \in \mathcal{K}^*. \end{array}$$

Assuming strong duality, we get that $c^T x + b^T y = 0$, so the dual and primal problems can be embedded into the following primal-dual feasibility problem

$$\begin{array}{ll} \text{find} & (x, s, y) \\ \text{subject to} & \begin{bmatrix} A & I & 0 \\ 0 & 0 & -A^T \\ c^T & 0 & b^T \end{bmatrix} \begin{bmatrix} x \\ s \\ y \end{bmatrix} = \begin{bmatrix} b \\ c \\ 0 \end{bmatrix} \\ & (s, y) \in \mathcal{K} \times \mathcal{K}^*. \end{array}$$

This is a feasibility problem with one affine subspace and one product of convex cones. There are many other ways to construct a feasibility problem with an affine subspace and a product of convex cones. One example is the homogeneous self-dual embedding which is often used in interior-point methods [18] and in the first-order optimization solver SCS [16]. Therefore, most convex optimization problems (at least those that can be posed as cone programs) can be solved using GAP, with one affine subspace and one product of convex cones. This is precisely the formulation for which the basic line search and the projected line search can be carried out with little additional cost and where the line search condition for the projected line search is convex in the line search parameter.

7 Numerical example

In this section we demonstrate the performance improvements of GAP when line search is used. We consider the following problem

$$\begin{array}{ll} \text{find} & z \\ \text{such that} & Q(z-p) = 0 \\ & z \geq 0, \end{array} \tag{24}$$

where $p = 10^{-7}\mathbf{1}$ to guarantee feasibility of the problem, and $Q \in \mathbb{R}^{50\times 100}$ is randomly generated with independent normally distributed elements with unit variance and zero mean.

We define two sets as $C = \{z \mid Q(z-p) = 0\}$ and $D = \{z \mid z \geq 0\}$. Depending on Q, the feasible set $C \cap D$ may be very small or consist of infinitely long rays in the nonnegative orthant. For this particular problem, Q is generated such that no ray in the affine set lies completely in the nonnegative orthant. Therefore, the intersection is relatively small.

As a termination criteria we use the following high accuracy requirement:

$$\begin{aligned} \|Q(z-p)\|_2 &\leq 10^{-10} \\ z &\geq 0, \end{aligned}$$

and we let $\alpha_1 = \alpha_2 \in [1, 2]$ and $\alpha = 0.85/\beta$, with β in (5).

As proposed in [11], we do not perform line search in each iteration, but use the rule

$$\frac{\langle r^k, \bar{r}^k \rangle}{\|r^k\|_2 \|\bar{r}^k\|_2} < 1 - 10^{-4} \tag{25}$$

to trigger it. The reason is that this often improves performance more than if line search is used in every iteration. The rationale behind the rule is that a large α_k can often be accepted when the iterates are moving along a straight line, i.e. when the angle between consecutive iterates is small. Numerical experiments suggest that consecutive iterates along a line seems to coincide with slow convergence, further motivating the use of line search when this occurs.

Both the basic and the projected line search is performed using a simple forward-tracking with a factor 1.4, and the result for different $\alpha_1 = \alpha_2$ are shown in Figure 2. The norm of the residual for each iteration is shown in Figure 3, for two different $\alpha_1 = \alpha_2$.

Without line search, it is clear from Figure 2 that the choice $\alpha_1 = \alpha_2 = 1$ and $\alpha_1 = \alpha_2 = 2$, corresponding to the Alternating Projections and Douglas-Rachford splitting respectively, are far from the optimal choice. They require approximately 5000 and $8 \cdot 10^6$ iterations compared to only 113 iterations for the optimal $\alpha_1 = \alpha_2$. However, this behavior does not apply to all problems, for example, in many cases the number of iterations seems to be monotonically decreasing with larger $\alpha_1 = \alpha_2$.

Figure 2 also reveals that the basic line search can considerably improve performance, especially for low choices of α_1, α_2, while the improvement for larger values is more modest. However, the projected line search performs considerably better, even for large $\alpha_1 = \alpha_2$, with only 52 iterations for the optimal $\alpha_1 = \alpha_2$. In particular, it decreases the iterations for $\alpha_1 = \alpha_2 = 2$ with more than a factor 10^5.

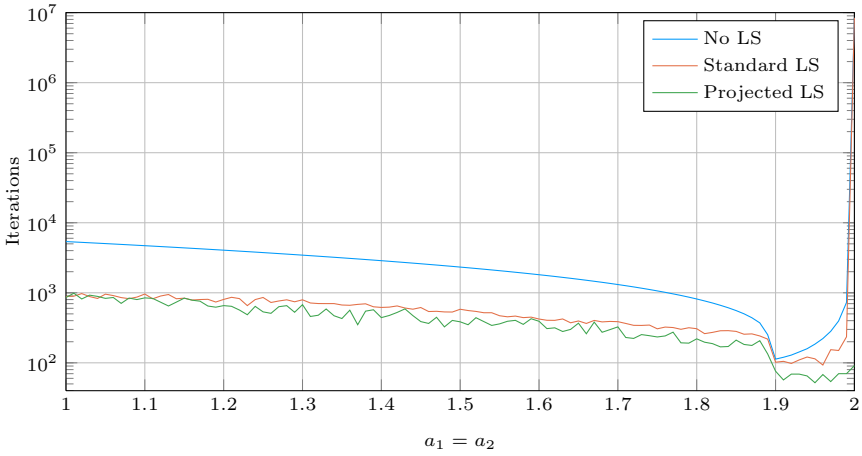

Figure 2: Number of iterations to solve problem (24) for different α_1, α_2, with and without line search.

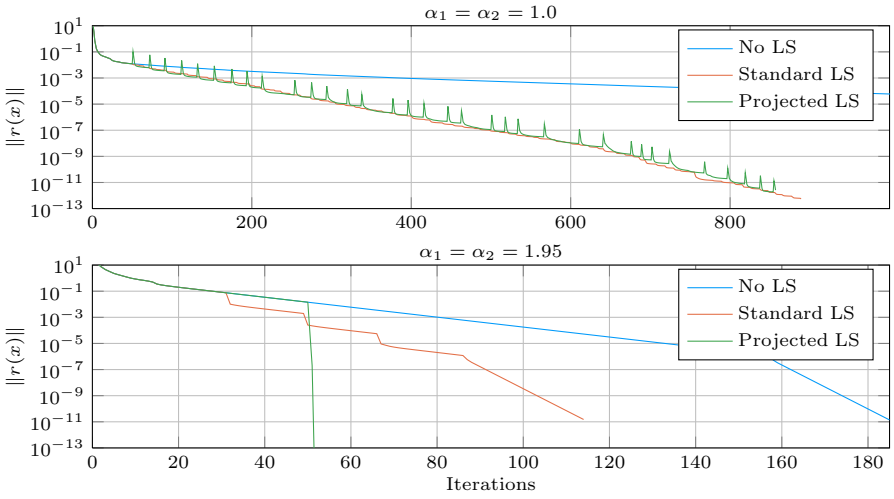

Figure 3: Norm of the residual for each iteration when solving problem (24) with different settings. It can be noted that the norm is strictly decreasing both without line search and with the standard line search. The peaks for the projected line search represents when a candidate α_k was accepted, which sometimes result in a temporary increase in the norm due to the constructed line search condition.

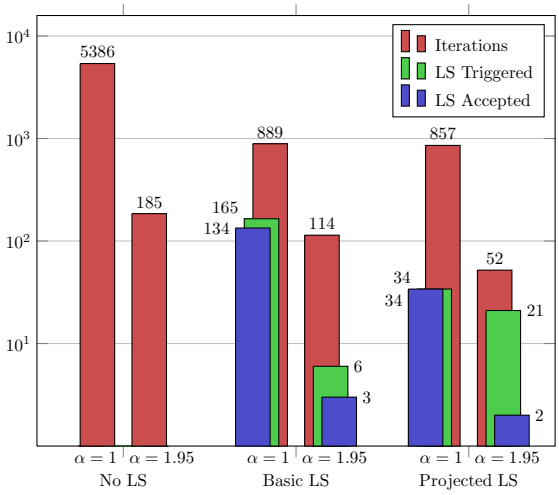

Figure 4: Number of times the line search was triggered and accepted for different algorithms and settings.

So far, we only compared the number of iterations for the different methods. Since the line search methods have a higher per iteration cost, we now evaluate what is actually gained by performing line search. We focus on the two cases with $\alpha_1 = \alpha_2 = 1.0$ and 1.95. Figure 4 shows the number of times the trigger criterion for line search (25) was satisfied for the standard and projected line search. It also shows how many times the line search found a point that satisfied the corresponding criterion for line search acceptance, i.e. (8) and (23). The number of evaluated candidate points (i.e. different α_k) averaged around 10 for each line search attempt, with a maximum of 18. Since $C = \{z \mid Q(z-p) = 0\}$ is affine, only one extra projection on C was needed for line search (to initialize the algorithm). To evaluate the acceptance criterion, (8) or (23), a few vector operations and one projection onto $D = \{z \mid z \geq 0\}$ is needed for each α_k. But projecting onto D is simply a max-operation and is thus very cheap.

8 Conclusions

We have shown that a recently proposed line search [11] is applicable to the generalized alternating projections method. We have also proposed an alternative line search method for GAP, the projected line search. Furthermore, we have shown that the line search condition for the projected line search is convex in the step length parameter. Both line search methods were evaluated on a feasibility problem, and showed significant performance improvements compared to the nominal method.

Acknowledgment

Both authors are financially supported by the Swedish Foundation for Strategic Research and members of the LCCC Linneaus Center at Lund University.

References

[1] S. Agmon. The relaxation method for linear inequalities. *Canadian Journal of Mathematics*, 6(3):382–392, 1954.

[2] H. H. Bauschke. The approximation of fixed points of compositions of nonexpansive mappings in hilbert space. *Journal of Mathematical Analysis and Applications*, 202(1):150 – 159, 1996.

[3] H. H. Bauschke and J. M. Borwein. On projection algorithms for solving convex feasibility problems. *SIAM Review*, 38(3):367–426, 1996.

[4] H. H. Bauschke and P. L. Combettes. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. Springer, 2011.

[5] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, New York, NY, 2004.

[6] L.M. Bregman. The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *USSR Computational Mathematics and Mathematical Physics*, 7(3):200 – 217, 1967.

[7] P. L. Combettes and I. Yamada. Compositions and convex combinations of averaged nonexpansive operators. *Journal of Mathematical Analysis and Applications*, 425(1):55–70, 2015.

[8] S. Diamond and S. Boyd. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 2016. To appear.

[9] J. Douglas and H. H. Rachford. On the numerical solution of heat conduction problems in two and three space variables. *Trans. Amer. Math. Soc.*, 82:421–439, 1956.

[10] P. Giselsson. Tight global linear convergence rate bounds for Douglas-Rachford splitting. 2015. Submitted. Available: `http://arxiv.org/abs/1506.01556`.

[11] P. Giselsson, M. Fält, and S. Boyd. Line search for averaged operator iteration. Available: `http://arxiv.org/abs/1603.06772v2`, 2016.

[12] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 3.0. `http://cvxr.com/cvx`, 2016.

[13] L. G. Gubin, B. T. Polyak, and E. V. Raik. The method of projections for finding the common point of convex sets. *USSR Computational Mathematics and Mathematical Physics*, 7(6):1–24, 1967.

[14] P. L. Lions and B. Mercier. Splitting algorithms for the sum of two nonlinear operators. *SIAM Journal on Numerical Analysis*, 16(6):964–979, 1979.

[15] J. Nocedal and S. Wright. *Numerical optimization*. Springer series in operations research and financial engineering. Springer, New York, NY, 2nd edition, 2006.

[16] B. O'Donoghue, E. Chu, N. Parikh, and S. Boyd. Conic optimization via operator splitting and homogeneous self-dual embedding. *Journal of Optimization Theory and Applications*, 2016.

[17] M. Udell, K. Mohan, D. Zeng, J Hong, S. Diamond, and S. Boyd. Convex optimization in Julia. *SC14 Workshop on High Performance Technical Computing in Dynamic Languages*, 2014.

[18] Y. Ye, M. J. Todd, and S. Mizuno. An o($\sqrt{n}l$)-iteration homogeneous and self-dual linear programming algorithm. *Mathematics of Operations Research*, 19(1):53–67, 1994.

A Appendix

A.1 Proof of Proposition 1

We start with the first claim. We know from [4, Proposition 4.8] that Π_{C_i} is firmly nonexpansive, and since $\alpha_i \in (0, 2)$ we know from [4, Corollary 4.29] that $P_{C_i}^{\alpha_i}$ are $\frac{\alpha_i}{2}$-averaged.

The composition $P_{C_p}^{\alpha_p} \dots P_{C_1}^{\alpha_1}$ is therefore β-averaged with β in (5), according to [7, 10]. Therefore

$$\begin{aligned} T &= (1-\alpha)\mathrm{Id} + \alpha((1-\beta)\mathrm{Id} + \beta \mathrm{S}) \\ &= (1-\alpha\beta)\mathrm{Id} + \alpha\beta \mathrm{S}. \end{aligned}$$

where S is nonexpansive. Since $\alpha \in (0, \frac{1}{\beta})$ we have $\alpha\beta \in (0, 1)$ and the first claim is proven.

The second claim is proven by noting that $P_{C_i}^{\alpha_i}$ is nonexpansive when $\alpha_i = 2$ [4, Corollary 4.10]. This implies that the composition is nonexpansive and the claim follows directly since $\alpha \in (0, 1)$.

A.2 Proof of Proposition 2

To show this, we need the following lemma.

Lemma 1 *Suppose that C is a nonempty closed and convex set and $\alpha \neq 0$. Then $\mathrm{fix} P_C^\alpha = C$, i.e. $x \in C$ if and only if $P_C^\alpha x = x$.*

Proof. It holds for projection operators with $\alpha = 1$ [4, Equation 4.8]. Since

$$P_C^\alpha x = \alpha \Pi_C x + (1-\alpha)x = x + \alpha(\Pi_C x - x)$$

we have $P_C^\alpha x = x$ if and only if $\Pi_C x = x$ if $\alpha \neq 0$. This concludes the proof of the lemma. □

The result follows directly for the case A1 from [4, Corollary 4.37] since $\mathrm{fix} P_{C_i}^{\alpha_i} = C_i$ and $P_{C_i}^{\alpha_i}$ are α_i-averaged operators.

For the case A2, let j be the index with $\alpha_j = 2$ and first assume that $j \neq 1, j \neq p$. Define $S_1 = P_{C_p} \dots P_{C_{j+1}}$ and $S_2 = P_{C_{j-1}} \dots P_{C_1}$.

Since all $P_{C_i}^{\alpha_i}$ are averaged for $i = j+1, \dots, p$, and since $\mathrm{fix} P_{C_i}^{\alpha_i} = C_i$ from Lemma 1, [4, Corollary 4.37] gives that S_1 is strictly quasi-nonexpansive and

that $\text{fix}S_1 = \cap_{i=j+1}^{p} C_i$. The same argument shows that S_2 is stricty quasi-nonexpansive with $\text{fix}S_2 = \cap_{i=1}^{j-1} C_i$.

Let $T_1 = S_1 P_{C_j}^2$. Nonexpansiveness of $P_{C_j}^2$ implies quasi-nonexpansiveness, so T_1 is also quasi-nonexpansive with $\text{fix}T_1 = \cap_{i=j}^{p} C_i$ by [4, Proposition 4.35]. Again applying [4, Proposition 4.35] to $T = T_1 S_2$ gives the result.

In the special case where $j = p$ or $j = 1$, the results follows in the same way for $T = P_{C_p}^2 S_2$ or $T = S_1 P_{C_1}^2$ respectively.

A.3 Proof of Theorem 2

For the new iterate x^{k+1} in (22b), we have $x^{k+1} = \Pi_C(x^k + \alpha_k r^k)$, and therefore $x^{k+1} \in C$. Let the shortest distance to a set Ω be denoted $\text{dist}_\Omega(\text{x}) := \|\Pi_\Omega \text{x} - \text{x}\|_2$. The norm of the residual then simplifies to

$$\begin{aligned}
\|r(x^{k+1})\| &= \|P_D^{\alpha_2} P_C^{\alpha_1} x^{k+1} - x^{k+1}\| && (26)\\
&= \|P_D^{\alpha_2} x^{k+1} - x^{k+1}\| && (27)\\
&= \|\alpha_2 \Pi_D x^{k+1} + (1-\alpha_2) x^{k+1} - x^{k+1}\| && (28)\\
&= \alpha_2 \|\Pi_D x^{k+1} - x^{k+1}\| && (29)\\
&= \alpha_2 \text{dist}_D(x^{k+1}) && (30)\\
&= \alpha_2 \text{dist}_D\left(\Pi_C(x^k + \alpha_k r^k)\right). && (31)
\end{aligned}$$

Since C is affine, so is Π_C. This implies that $\Pi_C(x^k + \alpha_k r^k)$ is affine in α_k. So the norm of the residual is the composition between the convex function dist_D and an affine function in α_k, hence convex [5] in α_k.